

REMARKS

In response to the Office Action dated March 28, 2003, Applicant respectfully requests reconsideration and withdrawal of the rejection of the claims and objection to the disclosure.

The Abstract was objected to because it contains more than 150 words. In response thereto, the Abstract has been amended to fall within the suggested word limit.

Claims 1-6 were rejected under 35 U.S.C. §102, on the grounds that they were considered to be anticipated by the Crouse patent (U.S. Patent No. 5,764,972). It is respectfully submitted, however, that the Crouse patent neither anticipates, nor otherwise suggests, the mirror file system of the present invention.

The mirror file system of the present invention is a virtual file system that mounts, links, and replicates two or more file systems on a single mount point, in real time. Once mounted and linked, the file systems and their stored components, i.e. directories and files, become mirroring partners. The mirror file system makes information available simultaneously on multiple network servers. However, the real-time mirroring mechanism is transparent to user applications. When the mirror file system receives updates from an application, it performs real-time updates on all file systems that are mounted by the MFS mounting protocol. The file systems can be local or exported by a remote network server. The mirror file system owns, manages and controls access to all the mounted file systems and their tree structures, enabling multiple network servers to own copies of the same information, which the mirror file system distributes and synchronizes in real time.

Unlike a conventional file system, the mirror file system of the present invention is a virtual, or stackable, file system that does not occupy physical storage. Rather, as illustrated in Figure 4, it is layered on top of two or more conventional file systems, e.g. local or network file systems. In essence, therefore, the mirror file system of the present invention can be viewed as a type of meta file system since it resides on top of, and controls, multiple conventional file systems.

It is respectfully submitted that the Crouse patent does not disclose an arrangement having the same architecture or functionality as the claimed invention. Rather, it is directed to a conventional file system that owns, controls and manages physical storage directly. However, it does not have any other file systems beneath it, nor otherwise under its control. Turning to the claims, claim 1 recites, among other elements, "means for mounting components of *each* of said *two* physical file systems in a single directory." With reference to this claimed subject matter, the Office Action refers to the Crouse patent at column 8, lines 12-17. This portion of the patent states that "the operating system program 20 will issue a mount command to the standard file control program 24 for a *particular* file tree 29, 38 in order to allow the user programs 22 at a particular user node 10 to have access to all of the files 26, 32 stored under that *particular* file tree 29, 38" (emphasis added). This portion of the patent merely describes the well-known fact that a file system has to be mounted in order for its file tree to be accessed. Applicants are not claiming the concept of mounting, per se. Rather, the claim recites that components of each of *two* physical file systems are mounted in a single directory. In contrast, the Crouse patent only

→ MERGED

} * CLAIM

discloses the concept of mounting a "particular" file tree, i.e. a *single* file system. It does not disclose that two file systems, or portions thereof, are mounted in a single directory.

Claims 1 goes on to recite a virtual file system data structure that contains elements which respectively correspond to each of the mounted components of the two file systems. The claim recites that each of these elements has "an application interface data structure with two associated pointers that respectively point to application interface data structures of a corresponding component in each of said two physical file systems." In the disclosed embodiment, the application interface data structure of the virtual file system elements comprises an mnode, which contains an associated vnode structure, and two vnode pointers. The two vnode pointers respectively point to the vnodes of the corresponding files or directories within the two file systems. In the example of Figure 4, these vnode pointers are represented in the mirror file system 203 as dark squares which respectively point to the corresponding directories or files in the file systems 201 and 202.

In connection with this claimed subject matter, the Office Action refers to the Crouse patent at column 11, lines 48-51. This portion of the patent states that the control information in a super block 70 includes a root directory pointer 86 that points to an inode 74 containing the logical address of the directory blocks that store the files and directories in the root directory for the file tree 38. The patent does not state, however, that each element of a virtual file system contains *two* pointers that respectively point to application interface data structures, e.g. vnodes, of corresponding components in each of *two* physical file systems.

Claim 2 recites that the application interface data structures correspond to a vnode structure. The rejection of this claim notes that the Crouse patent discloses a vnode at column 10, lines 44-54. Again, however, Applicant is not claiming the concept of a vnode, per se. Rather, the claims are directed to the manner in which a vnode structure is employed in the context of the present invention to enable an application to access a component stored within a conventional file system under the control of the mirror file system. It is respectfully submitted that the disclosure of a vnode in the Crouse patent does not suggest this claimed concept to one of ordinary skill in the art.


In summary, the claimed subject matter is directed to a mounting protocol and data structure that enables two or more different and separate physical file systems to be mounted on a single mount point, i.e. a directory. The data structure contains pointers and information about itself and about the two or more different and separate file systems, so that the mirror file system can control and manage these file systems beneath it. The mounting protocol and data structure provide complete transparency to the applications above the mirror file system, in user space, as well as the underlying file systems controlled by the mirror file system. In contrast, the archiving file system disclosed in the Crouse patent is a conventional file system. In particular, it mounts only one physical file system, belonging to one server system, on a single mount point. As such, it does not anticipate the subject matter of claim 1, nor its dependent claims.

The same distinctions are submitted to be applicable to claims 4-6 as well, and newly submitted claims 7-16. Reconsideration and withdrawal of the rejection, and allowance of all pending claims, are therefore respectfully requested.

Respectfully submitted,

BURNS, DOANE, SWECKER & MATHIS, L.L.P.

Date: June 30, 2003

By: 
James A. LaBarre
Registration No. 28,632

P.O. Box 1404
Alexandria, Virginia 22313-1404
(703) 836-6620